CASRE Three-Page Product Documentation

Allen Nikora, Jet Propulsion Laboratory, March 16, 2005

What it is: CASRE is a tool for estimating and forecasting software reliability based on the failure history of a software system during test. It accepts two types of failure history as input: time between successive failures, or the number of failures in a test interval and the length of a test interval. CASRE implements about a dozen of the most widely-used software reliability models, and operates under all versions of Windows. CASRE is available at no charge from the Open Channel Foundation at http://www.openchannelfoundation.org/projects/CASRE_3.0.

Features: CASRE displays its results in both graphical and tabular form. In addition to its modeling capabilities, CASRE provides 1) trend tests (running arithmetic average and Laplace test) to determine whether it is appropriate to apply software reliability models to a software development effort, 2) capabilities for identifying the most appropriate model (goodness of fit, prequential likelihood ratio, model bias, and model bias trend), and 3) the ability to combine the results of several models according to a static or dynamic weighting scheme to increase model accuracy.

As shown in Figure 1, CASRE has three windows for displaying data, data analyses, and modeling results. The window in the upper left shows the failure history data on which the models implemented in CASRE will operate. CASRE's models can operate on two types of failure data – times between successive failures ("interfailure times") or the number of failures observed in a test interval of given length ("interval data"). The failure data shown in Figure 1 is interval data – the columns represent the test interval sequence number, the number of failures observed in that test interval, the length of the interval, and the severity of those failures, on an arbitrary scale from 1 to 9. If failures of different severities were observed during the same interval, there will be as many lines displayed for that interval as the number of different severities observed. Filters can be applied to view only those failures within a specified severity range, or having a single severity value. The filters preserve the correct interfailure time or failures per interval values. Software developers often want to know things such as the expected time to the next failure, the failure intensity, and reliability in terms of the most critical failures only – the filtering capability allows users to easily make these computations by simply applying the appropriate severity filter before applying software reliability models. Figure 1 displays the failure data, trend analysis results, and model results. In Figure 1, the failure data is plotted using the "+" symbol; the estimates and predictions of the models applied to that data are shown using colored symbols. In this case, CASRE makes predictions of failure intensity for the next five intervals beyond the end of the failure data; this is shown by the five colored points for each model beyond the I-beam symbol denoting the end of the failure data set. CASRE can display the model results in many different forms, including cumulative number of failures, failure intensity, interfailure times, and reliability. CASRE also displays quantitative model evaluation criteria such as goodness of fit, prequential likelihood ratio (how much more likely is it that one model will produce accurate predictions than another?), model bias, and model bias trend. CASRE also allows the user to rank models according to these criteria.

Finally, the bottom window shown in Figure 1 displays model results and raw data in tabular form. The model's estimates and predictions, parameter values, and evaluation criteria are displayed in form that can be exported to a comma separated value file, which in turn can be brought into a spreadsheet or statistical program for further analysis. This feature is aimed more at software reliability researchers wishing to investigate the behavior of models in detail, although software development organizations can also make use of results in this form.

Benefits: CASRE helps software developers determine when to stop testing by answering these questions: 1) Does the software currently meet its reliability requirements? 2) How much more effort is required to achieve the required reliability? 3) When will no more failures be observed?

Successes: The development and maintenance of CASRE was funded by the U. S. Air Force Operational Test and Evaluation Center (AFOTEC) for use in evaluating the reliability of several software systems developed for the Air Force. It has also been used by commercial software development organizations – Sun Microsystems, the Hewlett Packard printer division, Raytheon, and Motorola have all used CASRE to evaluate their software systems. Over 400 copies of CASRE have been downloaded from the Open Channel Foundation. The organizations that downloaded CASRE include SAIC, Alcatel, Siemens, Northrop Grumman, and Cisco Systems – however, it is not possible to determine how it is being used within those organizations.

Contexts in which it is best used: CASRE ***must*** be used after unit test, during the software integration or system integration phases. In addition, the testing effort ***must*** be similar to the way in which the system is expected to be operated during fielded use – if the input space is divided into equivalence classes, the frequency with which those classes are sampled during testing must be similar to the frequency with which those classes are expected to be sampled during operational use. Otherwise, the system may not exhibit reliability growth during test, and the models would not produce meaningful results. A software development effort using CASRE must also be diligent in recording failures observed during test as well as the elapsed test time.

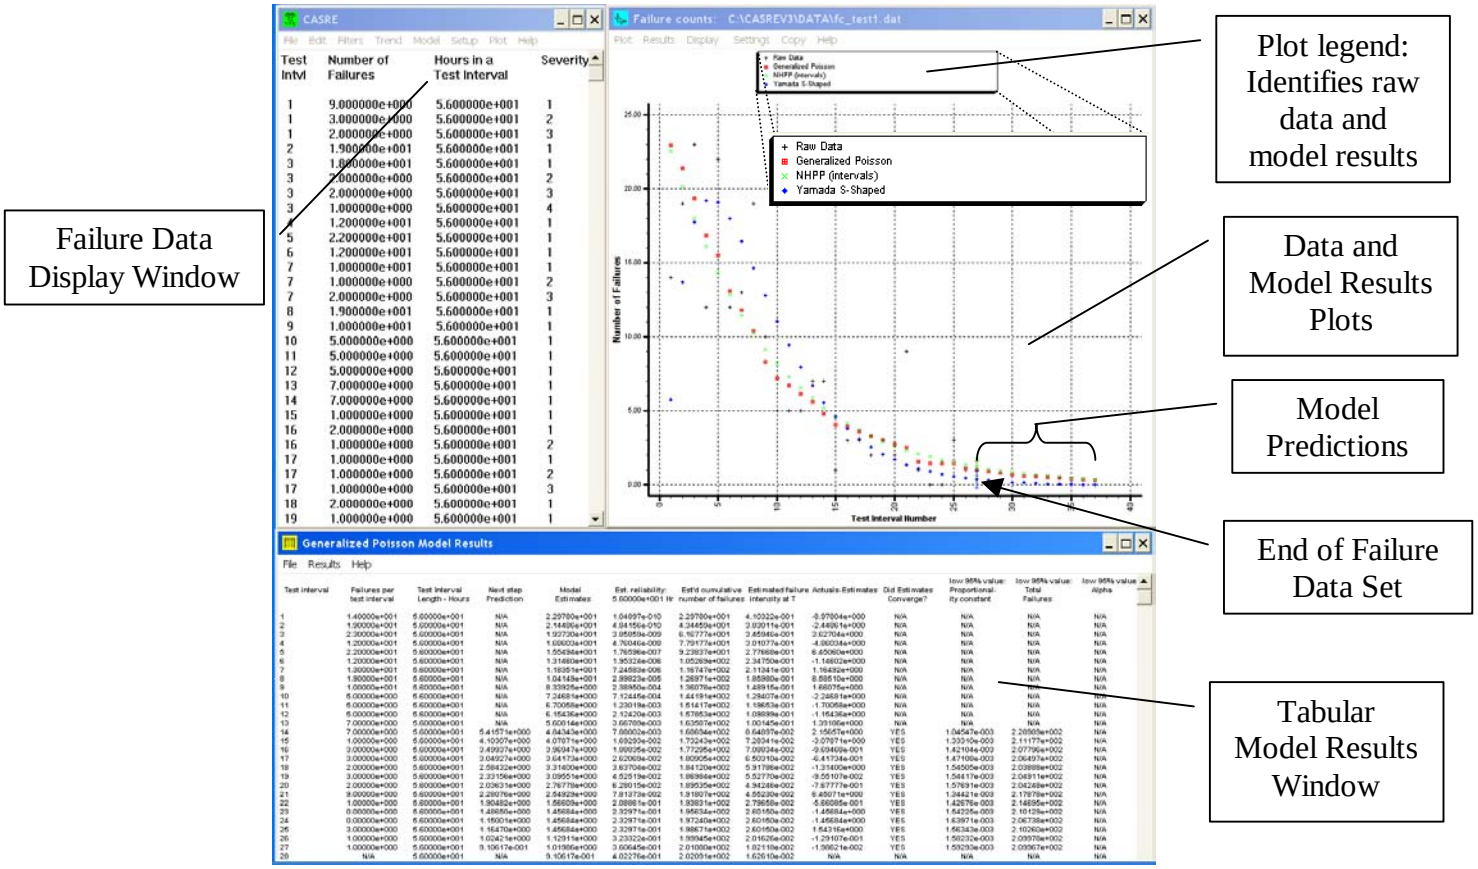

Figure 1 - CASRE On-Screen Appearance

Compare with alternative known products or technologies. There are several commercially and publicly available software reliability tools. These include SMERFS (Statistical Modeling and Estimation of Reliability Functions for Software), SoRel (developed by LAAS in France),

and the set of tools developed by Ann Marie Neufelder. The best-known is SMERFS, developed by Bill Farr at the Naval Surface Warfare Center (http://www.slingcode.com/smerfs/). Indeed, CASRE uses the mathematical modeling libraries that were developed for SMERFS. CASRE differs from the current version of SMERFS in that it 1) implements trend tests that can be applied to failure data to determine whether it is even appropriate to apply software reliability models, and 2) the ability to combine model results according to user-specified static or dynamic weighting schemes to improve model accuracy. CASRE also has more extensive documentation than SMERFS. CASRE implements more models and has more sophisticated methods for analyzing a model's goodness of fit and applicability than do Ann Marie Neufelder's tools.

What will a successful collaboration look like?

a. **What will you as the technology provider do?** We will work with you during proposal development on planning your collaboration. For a collaboration using CASRE, we will provide a 2-day training course at your site, assistance in applying CASRE to failure data and analyzing the results, and customer telephone support over the course of the collaboration. We will also put you in touch with other groups, NASA or otherwise, using CASRE.
b. **What should the development team do?** Prior to the collaboration, the NASA software development team should be about to start or be conducting software integration or system integration testing. The testing should involve repeated sampling of the input space (which can be accomplished by regression testing after a new build). The development team should also be diligent in recording elapsed testing time (time during which the system was actually running tests) as well as failures observed during testing in order to provide accurate input to CASRE. During the collaboration, take the training course; identify the builds to which CASRE will be applied; apply CASRE, collect data on its performance, and compare with existing V&V performance data.
c. **How will you, as technology provider, work together with the development team to ensure a successful collaboration?** During the proposal process, we will start by ensuring a good technical and project match with our technology. In addition to providing customer support for CASRE, we will track the collaboration's schedule, and at appropriate points proactively contact the Collaboration PoC, learn about problems or potential problems, obtain preliminary data, offer suggestions on how to proceed, and follow up to ensure that CASRE is being used effectively and the collaboration is achieving its success criteria.